

CLAIMS:

What is claimed is:

- 1 1. A method comprising:
2 determining at least one characteristic of a memory request; and
3 selectively leaving an accessed memory page open after a memory access based, at least
4 in part, on the at least one characteristic for the memory request, to balance memory access
5 latency and bandwidth of a subsequent memory request(s).
- 1 2. A method according to claim 1, wherein the at least one characteristic is determined
2 based, at least in part, on whether the memory request or a subset of memory requests are to a
3 single memory page or to more than one memory page.
- 1 3. A method according to claim 2, wherein the single memory page is left open after a
2 memory access if the memory request or the subset of memory requests is to the single memory
3 page.
- 1 4. A method according to claim 2, wherein the single memory page is closed after a
2 memory access if the memory request or the subset of memory requests is to more than one
3 memory page.
- 1 5. A method according to claim 1, wherein the determining at least one characteristic of the
2 memory request is determined based, at least in part, on a type of memory request expected to be
3 received.

1 6. A method according to claim 5, wherein the type of memory request is an instruction
2 memory request.

1 7. A method according to claim 6, wherein the instruction memory request results in a page
2 management indicator for leaving the memory page open after the memory access.

1 8. A method according to claim 5, wherein the type of memory request is a data memory
2 request.

1 9. A method according to claim 8, wherein the data memory request results in a page
2 management indicator for closing the memory page after the memory access.

1 10. A method according to claim 1, wherein the at least one characteristic is determined,
2 based at least in part, on an arbitration scheme.

1 11. A method according to claim 10, wherein the arbitration scheme is based, at least in part,
2 on a priority of a memory request.

1 12. A method according to claim 11, wherein the priority is based, at least in part, on
2 fairness.

1 13. A method according to claim 11, wherein the priority is based, at least in part, on quality
2 of service.

1 14. A method according to claim 1, wherein a memory controller receives the memory
2 request.

1 15. An apparatus comprising:
2 a plurality of memory pages; and
3 a memory controller, coupled with the plurality of memory pages, to analyze at least a
4 subset of received memory requests, to determine whether to selectively leave an accessed
5 memory page open after a memory access based, at least in part, on whether the memory
6 request(s) are to a single memory page or to more than one memory page.

1 16. An apparatus according to claim 15, the apparatus further comprising a memory to store
2 content, at least a subset of which is executable content; and
3 a control logic, coupled with the memory, to selectively execute at least a subset of the
4 executable content, to implement an instance of a memory controller.

1 17. An apparatus according to claim 15, wherein the plurality of memory pages is associated
2 with physical elements of synchronous dynamic random access memory.

1 18. An apparatus according to claim 15, wherein the determination to selectively leave an
2 accessed memory page open after a memory access is dynamic.

1 19. An apparatus according to claim 15, wherein a memory controller receives the at least
2 subset of memory requests.

1 20. A memory controller comprising:
2 a plurality of memory pages; and
3 a page manager, coupled with the plurality of memory pages, to selectively leave an
4 accessed memory page open after a memory access based, at least in part, on at least one
5 characteristic for a memory request.

1 21. A memory controller according to claim 20, the memory controller further comprising a
2 memory to store content, at least a subset of which is executable content; and
3 a control logic, coupled with the memory, to selectively execute at least a subset of the
4 executable content, to implement an instance of the page manager.

1 22. A memory controller according to claim 20, wherein the accessed memory page is
2 associated with elements of synchronous dynamic random access memory.

1 23. A system comprising:
2 volatile memory, associated with a plurality of memory pages; and
3 a page manager, coupled with the volatile memory, to selectively leave an accessed
4 memory page open after a memory access based, at least in part, on at least one characteristic for
5 a memory request.

1 24. A system according to claim 23, wherein the at least one characteristic for the memory
2 request is determined based, at least in part, on a type of memory request expected to be received
3 from an agent making a memory request.

1 25. A system according to claim 23, wherein the at least one characteristic for the memory
2 request or a subset of memory requests is determined based, at least in part, on whether the
3 memory request or the subset of memory requests are to a single memory page or to more than
4 one memory page.

1 26. A system according to claim 23, wherein the at least one characteristic is determined,
2 based at least in part, on an arbitration scheme.

1 27. A system according to claim 26, wherein the arbitration scheme is based, at least in part,
2 on a priority of a memory request.

1 28. A system according to claim 27, wherein the priority is based, at least in part, on fairness.

1 29. A system according to claim 27, wherein the priority is based, at least in part, on quality
2 of service.

1 30. A system according to claim 23, wherein the volatile memory is synchronous dynamic
2 random access memory.

1 31. A storage medium comprising content, which, when executed by a machine, causes the
2 machine to:
3 determine at least one characteristic of a memory request; and
4 selectively leave an accessed memory page open after a memory access based, at least in
5 part, on the at least one characteristic for the memory request, to balance memory access latency
6 and bandwidth of a subsequent memory request(s).

1 32. A storage medium according to claim 31, wherein the at least one characteristic for the
2 memory request is determined based, at least in part, on a type of memory request expected to be
3 made by an agent making a memory request.

1 33. A storage medium according to claim 31, wherein the at least one characteristic for the
2 memory request or a subset of memory requests is determined based, at least in part, on whether
3 the memory request or the subset of memory requests are to a single memory page or to more
4 than one memory page.

1 34. A storage medium according to claim 31, wherein the at least one characteristic is
2 determined, based at least in part, on an arbitration scheme.

1 35. A storage medium according to claim 34, wherein the arbitration scheme is based, at least
2 in part, on a priority of a memory request.

1 36. A storage medium according to claim 35, wherein the priority is based, at least in part, on
2 fairness.

1 37. A storage medium according to claim 35, wherein the priority is based, at least in part, on
2 quality of service.